REFERENCES

Andersen, P. K., Borgan, O., Gill, R. D., and Keiding, N. (1993), *Statistical Models Based on Counting Processes*, New York: Springer.

Ansell, J. I., and Phillips, M. J. (1989), "Practical Problems in the Statistical Analysis of Reliability Data" (with discussion), *Applied Statistics*, 38, 205–247.

Ascher, H., and Feingold, H. (1984), *Repairable Systems Reliability*, New York: Marcel Dekker.

Barlow, R. E., and Proschan, F. (1975), *Statistical Theory of Reliability and Life Testing*, New York: Holt, Rinehart and Winston.

Becker, R. A., Clark, L. A., and Lambert, D. (1998), "Events Defined by Duration and Severity, With an Application to Network Reliability" (with discussion), *Technometrics*, 40, 177–194.

Berman, M., and Turner, T. R. (1992), "Approximating Point Process Likelihoods With GLIM," *Applied Statistics*, 41, 31–38.

Cox, D. R. (1972), "Regression Models and Life Tables" (with discussion), *Journal of the Royal Statistical Society*, Ser. B, 34, 187–220.

Cox, D. R., and Lewis, P. A. W. (1966), *The Statistical Analysis of Series of Events*, London: Methuen.

Crowder, M. C., Kimber, A. C., Smith, R. L., and Sweeting, T. J. (1991), *Statistical Analysis of Reliability Data*, London: Chapman and Hall.

Dalal, S. R., and McIntosh, A. A. (1994), "When to Stop Testing for Large Software Systems With Changing Code," *IEEE Transactions on Software Engineering*, 20, 318–323.

Duchesne, T., and Lawless, J. (2000), "Alternative Time Scales and Failure Time Models," *Lifetime Data Analysis*, 6, 157–179.

Hoyland, A., and Rausand, M. (1994), *System Reliability Theory*, New York: Wiley.

Kalbfleisch, J. D., and Prentice, R. L. (1980), *The Statistical Analysis of Failure Time Data*, New York: Wiley.

Lawless, J. F. (1982), *Statistical Models and Methods For Lifetime Data*, New York: Wiley.

——— (1983), "Statistical Methods in Reliability" (with discussion), *Technometrics*, 25, 305–335.

——— (1998), "Statistical Analysis of Product Warranty Data," *International Statistical Review*, 66, 41–60.

Lawless, J. F., and Nadeau, C. (1995), "Some Simple Robust Methods for the Analysis of Recurrent Events," *Technometrics*, 37, 158–168.

Makis, V., and Jardine, A. K. S. (1992), "Computation of Optimal Policies in Replacement Models," *Institute of Mathematics and its Applications Journal of Mathematics Applied in Business and Industry*, 3, 169–175.

Meeker, W. Q., and Escobar, L. A. (1998), *Statistical Methods for Reliability Data*, New York: Wiley.

Nelson, W. (1982), *Applied Life Data Analysis*, New York: Wiley.

——— (1990), *Accelerated Testing*, New York: Wiley.

Olin, B. D., and Meeker, W. Q. (1996), "Applications of Statistical Methods to Nondestructive Evaluation" (with discussion), *Technometrics*, 38, 95–130.

Singpurwalla, N. D. (1995), "Survival in Dynamic Environments," *Statistical Science*, 10, 86–103.

Therneau, T. M. (1997), "Extending the Cox Model," in *Proceedings of the First Seattle Symposium on Biostatistics*, eds. D. Y. Lin and T. R. Fleming, New York: Springer, pp. 51–84.

Whitmore, G. A., Crowder, M. J., and Lawless, J. F. (1998), "Failure Inference from a Marker Process Based on a Bivariate Wiener Model," *Lifetime Data Analysis*, 4, 229–251.

The State of Statistical Process Control as We Proceed into the 21st Century

Zachary G. Stoumbos, Marion R. Reynolds, Jr., Thomas P. Ryan, and William H. Woodall

1. INTRODUCTION

Statistical process control (SPC) refers to some statistical methods used extensively to monitor and improve the quality and productivity of manufacturing processes and service operations. SPC primarily involves the implementation of *control charts*, which are used to detect any change in a process that may affect the quality of the output. Control charts are among the most important and widely used tools in statistics. Their applications have now moved far beyond manufacturing into engineering, environmental science, biology, genetics, epidemiology, medicine, finance, and even law enforcement and athletics (see Lai 1995; Montgomery 1997; Ryan 2000). C. R. Rao (1989) stated: "It is not surprising that a recent book on modern inventions lists statistical quality control as one of the technological inventions of the past century. Indeed, there has rarely been a technological invention like statistical quality control, which is so wide in its application yet so simple in theory, which is so effective in its results yet so easy to adopt and which yields so high a return yet needs so low an investment."

The first control charts were developed by Walter A. Shewhart in the 1920s (see Shewhart 1931). These simple Shewhart charts have dominated applications to date. Much research has been done on control charts over the last 50 years, but the diffusion of this research to applications has been very slow. As Crowder, Hawkins, Reynolds, and Yashchin (1997) noted, "There are few areas of statistical application with a wider gap between methodological development and application than is seen in SPC."

An examination of what is used in practice and what appears in the SPC literature shows that there are actually two gaps. There is one gap between applications and applied re-

Zachary G. Stoumbos is Assistant Professor, Department of Management Science and Information Systems and Rutgers Center for Operations Research, Rutgers, The State University of New Jersey, Newark, NJ 07102 (E-mail: *stoumbos@andromeda.rutgers.edu*). Marion R. Reynolds, Jr. is Professor, Departments of Statistics and Forestry, Virginia Polytechnic Institute and State University, Blacksburg, VA 24061. Thomas P. Ryan is Visiting Professor, Department of Statistics, University of Michigan, Ann Arbor, MI 48109. William H. Woodall is Professor, Department of Statistics, Virginia Polytechnic Institute and State University, Blacksburg, VA 24061. The authors wish to express appreciation to the following individuals for their comments on the manuscript: J. Stuart Hunter, Princeton University; C. R. Rao, Penn State University; Chris P. Tsokos, University of South Florida; and Emmanuel Yashchin, IBM Research. Zachary G. Stoumbos' work was supported in part by 1999 and 2000 Rutgers Faculty of Management Research Fellowships. William H. Woodall's work was supported in part by National Science Foundation grant DMI 9908013.

Journal of the American Statistical Association
September 2000, Vol. 95, No. 451, Vignette

search in journals such as the *Journal of Quality Technology* and *Technometrics* and another gap between this applied research and the research in some of the more theoretical statistics journals. The existence of these gaps is disturbing, because it means that most practitioners have received little of the potential benefit from the technical advances made in SPC over the last half-century. Here we discuss in detail the current state of SPC and give our views on some important future research topics. These topics have good potential to narrow the gaps between applications and applied and theoretical SPC research.

2. THE PROCESS-MONITORING PROBLEM

The process-monitoring problem can be described in general terms as follows. Let X represent a quality variable of interest, and suppose that $f_{\boldsymbol{\theta}}(x)$, the distribution function of X, is indexed by $\boldsymbol{\theta}$, a vector of one or more parameters. A stable process that is operating with $\boldsymbol{\theta} = \boldsymbol{\theta}_0$ is said to be in *statistical control*. The value of $\boldsymbol{\theta}_0$ may or may not correspond to an ideal (or target) value.

"Murphy's law" explains the purpose of process monitoring; over time, something will inevitably change and possibly cause deterioration in process quality. Something that affects process quality is assumed to be reflected by a change in $\boldsymbol{\theta}$ from the value $\boldsymbol{\theta}_0$, so the basic goal of process monitoring is to detect changes in $\boldsymbol{\theta}$ that can occur at unknown times. Many types of changes in $\boldsymbol{\theta}$ could occur, such as brief self-correcting changes or shifts and drifts that persist over long periods if undetected.

Control charts for monitoring $\boldsymbol{\theta}$ are based on taking samples from the process and observing the values of X. A *control statistic*, say Y, is computed after each sample and plotted in time order on a control chart. *Control limits* are constructed such that a value of Y is very unlikely to fall outside of them when $\boldsymbol{\theta} = \boldsymbol{\theta}_0$. A value of Y that falls outside the control limits is taken as a signal that a change in $\boldsymbol{\theta}$ has occurred, and that some appropriate action is required.

At the start of the process, the values of some or all of the components of $\boldsymbol{\theta}_0$ may be unknown, and thus a preliminary phase of collecting process data, estimating parameters, and testing for process stability may be required. Process monitoring can begin after $\boldsymbol{\theta}_0$ is estimated in this preliminary phase.

The crisis and subsequent quality revolution in U.S. industry in the 1980s triggered an increasing emphasis on actively working to improve quality (Deming 1986). Thus, in addition to detecting undesirable changes in $\boldsymbol{\theta}$, control charts also should be used to identify improvements in the process. For example, if X is normally distributed and $\boldsymbol{\theta} = (\mu, \sigma)$, then improving quality might correspond to making process adjustments that will reduce σ (Reynolds and Stoumbos 2000a).

3. TRADITIONAL SHEWHART CONTROL CHARTS

The first control charts proposed by W. A. Shewhart in the 1920s remain in widespread use today. The Shewhart charts were designed to make it relatively easy for process personnel without statistical training to set up, apply, and interpret the charts using only a pencil and paper for calculations. Although it is not often explicitly stated, these charts are based on the assumption that $f_{\boldsymbol{\theta}}(x)$ is one of a few standard distributions (normal for continuous data, and binomial or Poisson for discrete data), and that successive observations of X are independent. The control statistic Y_k computed after sample k is a function of the data in sample k only. Ease of computation is emphasized, so, for example, the sample range is typically used as the measure of dispersion. The design of Shewhart charts is traditionally based on simple heuristics, such as using samples of four or five observations at suitable sampling intervals, say every hour, and using "three-sigma" limits set three standard deviations away from the in-control mean of Y_k.

Shewhart charts have functioned as simple graphical tools in a wide variety of applications. It is not surprising, however, that such simple charts are usually far from optimal (see Secs. 4 and 5) and may even be inappropriate. For example, with three-sigma limits, the false-alarm rate is not adjusted to suit the specific conditions of different applications, and anomalies arise such as lack of lower control limits for nonnegative statistics. The use of the sample range has continued long after computational ease ceased to be a primary concern.

Modifying a Shewhart chart can alleviate some of the aforementioned problems, but a greater disadvantage is that these charts are inefficient for detecting all but relatively large changes in $\boldsymbol{\theta}$. The increasing emphasis today on high-quality products increases the importance of detecting small changes in $\boldsymbol{\theta}$.

4. MORE EFFICIENT CONTROL PROCEDURES

Efficient detection of small and moderate shifts in $\boldsymbol{\theta}$ requires that the control statistic in some way accumulate information across past samples. *Runs rules*, which are based on patterns of points in a Shewhart chart, help to improve the ability of Shewhart charts to detect small shifts in $\boldsymbol{\theta}$ (Champ and Woodall 1987), but using these rules is not the best method of detecting small shifts in $\boldsymbol{\theta}$.

A much better method of accumulating information across samples uses a control statistic that is an *exponentially weighted moving average* (EWMA) of current and past sample statistics. In particular, if I_k is the individual statistic for sample k, then the EWMA control statistic computed after sample k is

$$E_k = (1-\lambda)^k E_0 + \sum_{i=1}^{k} (1-\lambda)^{k-1} \lambda I_i = (1-\lambda) E_{k-1} + \lambda I_k,$$

where E_0 is the starting value and $\lambda > 0$ is the smoothing parameter that determines the weight given to current data relative to past data. A signal is given if E_k falls outside of control limits.

The *cumulative sum* (CUSUM) chart is another highly efficient control chart that accumulates information over current and past samples. The CUSUM statistic for detecting a shift from $\boldsymbol{\theta}_0$ to a specified alternative $\boldsymbol{\theta}_1$ can be written

as

$$C_k = \max\{0, C_{k-1}\} + \ln(f_{\theta_1}(x_k)/f_{\theta_0}(x_k)),$$

where C_0 is the starting value. A signal is given if C_k exceeds a control limit. Two-sided CUSUM charts are usually constructed by running two one-sided CUSUM charts simultaneously.

For a given false-alarm rate, both the EWMA and CUSUM charts are much better than a Shewhart chart for detecting small sustained shifts in θ. The CUSUM chart is optimal for detecting a shift from θ_0 to a specified θ_1 in that it minimizes the worst mean signal delay for a large class of signal rules with appropriately constrained false-alarm rates (Lorden 1971; Moustakides 1986; Ritov 1990). A good broad exposition on CUSUM charts from a more applied perspective was given by Hawkins and Olwell (1998), and discussion of some recent developments on CUSUM charts have been provided by Reynolds and Stoumbos (1999, 2000b). For the problem of monitoring the process mean μ, several studies have shown that the EWMA and CUSUM charts generally have similar detection efficiencies over a range of shifts in μ (see, e.g., Lucas and Saccucci 1990). The EWMA and CUSUM charts date from the 1950s (Page 1954 for the CUSUM and Roberts 1959 for the EWMA), but usage of these efficient charts in applications was very infrequent for many years. Their usage is steadily increasing, although still relatively low.

A number of *generalized* CUSUM schemes proposed over the last three decades allow for θ_1 to be any unknown value in a given interval. These schemes are based on a generalized likelihood ratio (GLR) or on integrating the likelihood ratio with respect to a probability distribution of θ (Basseville and Nikiforov 1993; Lai 1995; references therein). The latter are quasi-Bayesian schemes akin to the procedures discussed in Section 5. Most generalized CUSUM schemes proposed to date are of mainly theoretical interest, because they cannot be easily expressed by computationally convenient recursive forms (Lai 1995).

5. BAYESIAN PROCEDURES AND ECONOMIC MODELS

Bayesian procedures appear to be naturally suited for process monitoring. The application of Bayesian procedures requires recognition that the form of the a priori information about θ is not simply a prior distribution, as would be the case in traditional estimation. In process monitoring, it is assumed that θ will eventually change from the value θ_0 while monitoring is conducted. Thus the prior distribution must reflect when a change in θ will occur, as well as the type of change that will occur.

Bayesian methods for process monitoring have been available since the works of Girshick and Rubin (1952), Shiryaev (1963), and Roberts (1966), who placed a geometric prior distribution with parameter p on the unknown time T of the change in θ (*changepoint*), and independently derived what is commonly termed the *Shiryaev–Roberts* (S-R) procedure. The S-R control statistic for detecting a shift from θ_0 to a specified alternative θ_1 can be expressed in terms of the log-likelihood ratio as

$$R_k = \ln(p + e^{R_{k-1}}) - \ln(1-p) + \ln(f_{\theta_1}(x_k)/f_{\theta_0}(x_k)),$$

where R_0 is the starting value. A signal is generated if R_k exceeds a control limit, which in the original work was determined to minimize a cost-based loss function.

Pollak (1985) proved the S-R signal rule to be asymptotically Bayes risk efficient as $p \to 0$. Basseville and Nikiforov (1993) and Lai (1995) provided good discussions of the S-R control chart. Bayesian procedures for process monitoring appear to be unknown to most applied statisticians and industrial engineers and thus are rarely used in SPC applications.

There is a relatively large volume of applied SPC literature on economic models (starting with Duncan 1956) related to Bayesian approaches. These economic models aim to find the optimum control chart design (sample size, sampling interval, and control limits) to minimize long-term expected costs. The control statistics used in these models are from standard control charts, such as Shewhart or CUSUM charts, and thus are not based on a posterior distribution. These models use prior distributions for the time and size of change in θ and a loss function that accounts for the costs of sampling, false alarms, and operating out of control. Thus these models have the key elements that would be used in a Bayesian approach to the problem, except that the control statistics are not based on a posterior distribution.

Economic models appear to provide a natural approach for process engineers to use in control chart design and application, because decisions are put into terms that managers understand (dollars), and the problem is framed in terms of designing traditional control charts. But, like many purely Bayesian methods, these models are rarely used by SPC practitioners. There is disagreement among SPC researchers about the general usefulness of the economic modeling approach. Some researchers (e.g., Woodall 1986) criticize these models and feel that future SPC research efforts would be more fruitful in other areas. Other researchers are actively working on these models and feel that they provide the best approach to control chart design for many applications (e.g., Keats, Del Castillo, von Collani, and Saniga 1997).

6. MORE EFFICIENT SAMPLING

The standard approach to sampling for a control chart is to use a *fixed sampling rate* (FSR) in which samples of fixed size are obtained using a fixed-length sampling interval. In recent years, *variable sampling rate* (VSR) control charts have been developed. VSR charts allow the sampling rate to vary as a function of the process data. When the data exhibit no evidence of a change in θ, a low sampling rate is used, but as soon as there is evidence of a possible change in θ, a high sampling rate is used. If the evidence of a change in θ is strong enough, a VSR chart signals in the same way as a traditional FSR chart. Using a high sampling rate when

there is evidence of a change in θ results in much faster detection of most shifts in θ, compared to an FSR chart with the same average in-control sampling rate.

There are several ways to allow the sampling rate to vary as a function of the process data. One way is to allow the sampling interval to vary (Reynolds, Amin, and Arnold 1990; Reynolds and Stoumbos 2000a; references therein). Another way is to allow the sample size to vary. A particularly efficient way to allow the sample size to vary is to apply a *sequential probability ratio test* for testing θ_0 versus θ_1 at each sampling point (Reynolds and Stoumbos 1998; Stoumbos and Reynolds 1997, 2000a). Tagaras (1998) presented a review of VSR charts.

The great majority of VSR control charts in the literature, including those mentioned herein, have been developed for discrete-time models. Recently, several VSR control charts have been developed for the continuous-time problem of monitoring the drift coefficient of a Brownian motion process (Assaf, Pollak, and Ritov 1992; Assaf and Ritov 1989; Srivastava and Wu 1994). Assaf et al. (1992) and Srivastava and Wu (1994) noted that these continuous-time VSR control charts are quite complicated to implement in practice and were considered mainly from a theoretical viewpoint, using diffusion theory.

The disadvantage of VSR charts is, of course, the administrative inconvenience of the varying sampling rate. However, the ability to make better use of sampling resources by selectively allocating them to the time periods in which they will be most effective provides a powerful method for significantly increasing the efficiency of process monitoring.

7. MULTIVARIATE CONTROL CHARTS

In many SPC applications, the quality of the process will be characterized by multiple correlated quality variables, and in this situation both the quality characteristic $\mathbf{X}$ and parameter $\boldsymbol{\theta}$ will be vectors. In multivariate SPC applications, the most common approach to process monitoring is to apply separate univariate control charts for each variable, ignoring the issue of their joint performance.

One approach to constructing a multivariate control chart is based on forming a single control statistic from the multivariate data in each sample. This control statistic would usually be a quadratic form involving summary statistics for each variable, and would be plotted on a Shewhart-type chart (Hotelling 1947). The resulting control chart has the disadvantage of all Shewhart-type charts; it is inefficient for detecting small and moderate-sized sustained shifts in $\boldsymbol{\theta}$.

A much better approach is to compute an EWMA or CUSUM statistic for each variable, and then use a quadratic form to combine these separate univariate statistics into a single control statistic to be plotted on a control chart as usual (Lowry, Woodall, Champ, and Rigdon 1992; Mason, Champ, Tracy, Wierda, and Young 1997; references therein).

8. AUTOCORRELATION

A basic assumption usually made in constructing and evaluating control charts is that the process data are independent. But autocorrelation is present in many applications, particularly in cases in which data are closely spaced in time. Relatively low levels of autocorrelation can have a significant impact on the statistical properties of standard control charts designed under the assumption of independence. For example, estimates of θ_0 can be severely biased, resulting in a much higher false-alarm rate than expected. It is not uncommon in applications for standard control charts to be applied to autocorrelated data. When these control charts do not work properly, ad hoc adjustments are made to the charts to try to compensate for the autocorrelation. This clearly is not the best approach to use for the problem of autocorrelation.

In recent years, two basic approaches to dealing with autocorrelation have been studied in the applied SPC literature. Under both approaches, an underlying time series model is assumed. The first approach uses the original data in a standard control chart, but adjusts the control limits to account for the autocorrelation. The second approach advocates plotting the residuals from the time series model on a standard control chart (Faltin, Mastrangelo, Runger, and Ryan 1997; Lu and Reynolds 1999; references therein).

Both of these approaches tend to make the problem of process monitoring appear simpler than it actually is. If a time series model really captures the in-control behavior of the process, then the parameters of this time series model become elements of $\boldsymbol{\theta}$. Thus the complexity of the process-monitoring problem is increased due to the increase in the number of parameters to estimate and monitor.

Several interesting extensions of CUSUM, GLR, and nonlikelihood control chart schemes for autocorrelated data have appeared in the engineering literature over the last three decades. Basseville and Nikiforov (1993) gave a comprehensive overview of these algorithms in the context of univariate as well as multivariate process-monitoring applications.

9. STATISTICAL PROCESS CONTROL AND AUTOMATIC PROCESS CONTROL

The basic philosophy of SPC for improving quality is to detect process changes, so that the cause(s) of the changes can be investigated. Another approach to improving quality, sometimes called *automatic process control* (APC), has been developed in the engineering literature. APC can be used in situations in which there is autocorrelation in the data and a mechanism is available to adjust the process when it appears to be deviating from the desired state. The approach used in APC is to forecast the next observation, and then use the adjustment mechanism to adjust the process so that the observation will be closer to the desired state. In APC the process is assumed to be wandering in some sense, and the adjustment mechanism compensates for this wandering. Thus the basic philosophy of APC is to compensate for undesirable process changes, rather than to detect and remove them as in SPC.

The determination of the best adjustment to make in APC requires a model for process behavior. The optimal adjustment chosen for this model may not work well in the presence of a process change that alters the model. The need to

detect changes in the underlying model suggests combining SPC monitoring with APC adjustment to exploit their individual strengths. There has been recent work on combining SPC and APC, but more work is needed (Box and Luceño 1997; Tsung, Shi, and Wu 1999; references therein).

10. DIAGNOSTICS

Monitoring, in either the case of a time series model or the case of multiple variables, involves monitoring multiple parameters. When a control chart signals in this situation, the parameter(s) that have changed may be difficult to determine. In addition, with small parameter changes, pinpointing the time of the change may also be difficult. Thus an important issue with multiple parameters is the ability to diagnose the type and time of the parameter change that has occurred.

In the case of multivariate data, some work has been done on diagnostics, but little has been done in the case of time series models. The literature on estimating changepoints in a sequence of observations is substantial (Basseville and Nikiforov 1993 and references therein). But the problem of estimating the changepoint in process monitoring is different from the problem for a fixed sequence of observations because in process monitoring the estimation is done only after a signal by the control chart (Nishina 1992; Nishina and Peng-Hsiung 1996; references therein). The problem of diagnostics is one area where the use of Bayesian models, nonlinear filtering theory, and stochastic calculus seem quite natural and may prove very useful (Lai 1995; Stoumbos 1999; Yashchin 1997).

11. PARAMETER ESTIMATION AND NONPARAMETRIC PROCEDURES

Control chart performance is very sensitive to errors in estimating θ_0. For example, the false-alarm rate may be much higher or lower than expected unless the dataset used in the initial phase of estimating θ_0 is quite large. The situation is even worse in more complex situations when multiple parameters must be estimated (Adams and Tseng 1998; Lu and Reynolds 1999). The effect of errors in estimating θ_0 in complex models awaits additional study, and methods for compensating for these effects remain to be developed.

Traditional Shewhart-type charts are usually based on the assumption that if $f_\theta(x)$ is continuous, then it will be normal. Almost all work on multivariate control charts is based on the assumption that $f_\theta(x)$ is multivariate normal. In some cases, the central limit theorem can be used to justify approximate normality when monitoring means, but in numerous cases normality is an untenable assumption and one is unwilling to use another parametric model. A number of nonparametric methods are available for use in these cases. A more prominent role is expected for nonparametric methods. As data availability increases, nonparametric methods seem especially useful in multivariate applications where most methods proposed thus far rely on normality. Nonparametric multivariate control charts have been studied only very recently and much more research is needed (e.g., Liu 1995; Stoumbos and Jones 2000; Stoumbos, Jones, Woodall, and Reynolds 2000).

12. FUTURE DIRECTIONS

Advances in automated manufacturing systems coupled with advances in sensing and automatic inspection technology will continue to increase the volume of data available for drawing inferences about many processes. In some applications, this will change the inference problem from one dependent on scarce data to one based on plentiful data. However, the emphasis on higher quality will often require measuring more variables, which in some cases may be expensive and/or time-consuming. Thus we foresee no reduction in the need for efficient procedures for process monitoring.

In the future, we expect problems to be more diverse, with specialized monitoring methods required. Multiple quality variables, along with possible autocorrelation in these variables, will require more complex models with a large increase in the number of parameters to monitor.

The increasing complexity of problems encountered provides an opportunity to narrow the gaps between applications and applied and theoretical SPC research. The Shewhart charts that have dominated industrial applications over the past 75 years were designed to be extremely simple, with a one-size-fits-all approach to their design and implementation. In the case of relatively simple problems, arguments that CUSUM or EWMA charts have much better statistical properties convinced only some industrial practitioners to move beyond using the familiar Shewhart charts. As problems become more complex, the need for more sophisticated monitoring procedures will become critical and more obvious to all. Theoretical and applied research that addresses this need can have a major impact on applications.

The following are some additional research areas that we feel have very good potential for impacting applications (for some other research topics, see Montgomery and Woodall 1997 and Woodall and Montgomery 1999):

- Basic and applied research is needed on methods for monitoring multiple parameters that arise in models for the cases of single or multiple process variables and/or autocorrelation.
- When multiple parameters are to be monitored, methods are needed for diagnosing both the changepoint (Nishina 1992; Nishina and Peng-Hsiung 1996) and the specific parameter or parameters that have changed (Reynolds and Stoumbos 2000a). The application of monitoring procedures to complex problems may require sophisticated efforts in model fitting and parameter estimation, and the effects of the fitting and estimation on the procedures awaits much more study. We foresee that VSR approaches can significantly improve the effectiveness of change-point estimation.
- The robustness of fitted process models and monitoring procedures to model misspecification needs further study (Stoumbos and Reynolds 2000b). Nonparametric procedures for multivariate problems is an open field with great potential.
- Additional basic and applied research is needed on procedures that integrate SPC and APC methodology.

- The scope of SPC methods should be expanded to include the study of all variation over time throughout the entire production process (tracking), which usually includes numerous stages of process steps and product measurement (Agrawal, Lawless, and Mackay 1999; Hawkins 1991; Lawless, Mackay, and Robinson 1999).
- A greater synthesis of the theoretical changepoint and applied SPC literatures is very desirable. Moreover, excellent opportunities exist for cross-fertilization of ideas from other areas of statistics and stochastic modeling, including epidemiology, outlier detection, and especially stochastic calculus and financial mathematics (Stoumbos 1999).
- The use of complicated models will place even greater emphasis on the development of software. In many cases, software will need to be customized for particular applications.

REFERENCES

Adams, B. M., and Tseng, I. T. (1998), "Robustness of Forecast-Based Monitoring Schemes," *Journal of Quality Technology*, 30, 328–339.

Agrawal, R., Lawless, J. F., and Mackay, R. J. (1999), "Analysis of Variation Transmission in Manufacturing Processes—Part II," *Journal of Quality Technology*, 31, 143–154.

Assaf, D., Pollak, M., and Ritov, Y. (1992), "A New Look at Warning and Action Lines of Surveillance Schemes," *Journal of the American Statistical Association*, 87, 889–895.

Assaf, D., and Ritov, Y. (1989), "Dynamic Sampling Procedures for Detecting a Change in the Drift of a Brownian Motion: A Non-Bayesian Model," *The Annals of Statistics*, 17, 793–800.

Basseville, M., and Nikiforov, I. V. (1993), *Detection of Abrupt Changes: Theory and Applications*, Engelwood Cliffs, NJ: Prentice-Hall.

Box, G. E. P., and Luceño, A. (1997), *Statistical Control by Monitoring and Feedback Adjustment*, New York: Wiley.

Champ, C. W., and Woodall, W. H. (1987), "Exact Results for Shewhart Control Charts With Supplementary Runs Rules," *Technometrics*, 29, 393–399.

Crowder, S. V., Hawkins, D. M., Reynolds, M. R., Jr., and Yashchin, E. (1997), "Process Control and Statistical Inference," *Journal of Quality Technology*, 29, 134–139.

Deming, W. E. (1986), *Out of the Crisis*, Cambridge, MA: Center for Advanced Engineering Studies, Massachusetts Institute of Technology.

Duncan, A. J. (1956), "The Economic Design of $\bar{X}$ Charts Used to Maintain Current Control of a Process," *Journal of the American Statistical Association*, 51, 228–242.

Faltin, F. W., Mastrangelo, C. M., Runger, G. C., and Ryan, T. P. (1997), "Considerations in the Monitoring of Autocorrelated and Independent Data," *Journal of Quality Technology*, 29, 131–133.

Girshick, M. A., and Rubin, H. (1952), "A Bayes Approach to a Quality Control Model," *The Annals of Mathematical Statistics*, 23, 114–125.

Hawkins, D. M. (1991), "Multivariate Quality Control Based on Regression-Adjusted Variables," *Technometrics*, 33, 61–75.

Hawkins, D. M., and Olwell, D. H. (1998), *Cumulative Sum Control Charts and Charting for Quality Improvement*, New York: Springer-Verlag.

Hotelling, H. (1947), "Multivariate Quality Control," in *Techniques of Statistical Analysis*, eds. C. Eisenhart, M. W. Hastay, and W. A. Wallis, New York: McGraw-Hill.

Keats, J. B., Del Castillo, E., von Collani, E., and Saniga, E. (1997), "Economic Modeling for Statistical Process Control," *Journal of Quality Technology*, 29, 144–147.

Lai, T. L. (1995), "Sequential Changepoint Detection in Quality Control and Dynamical Systems" (with discussion), *Journal of the Royal Statistical Society*, Ser. B, 57, 613–658.

Lawless, J. F., Mackay, R. J., and Robinson, J. A. (1999), "Analysis of Variation Transmission in Manufacturing Processes—Part I," *Journal of Quality Technology*, 31, 131–142.

Liu, R. Y. (1995), "Control Charts for Multivariate Processes," *Journal of the American Statistical Association*, 90, 1380–1387.

Lorden, G. (1971), "Procedures for Reacting to a Change in Distribution," *The Annals of Mathematical Statistics*, 42, 1897–1908.

Lowry, C. A., Woodall, W. H., Champ, C. W., and Rigdon, S. E. (1992), "A Multivariate Exponentially Weighted Moving Average Control Chart," *Technometrics*, 34, 46–53.

Lu, C. W., and Reynolds, M. R., Jr. (1999), "EWMA Control Charts for Monitoring the Mean of Autocorrelated Processes," *Journal of Quality Technology*, 31, 166–188.

Lucas, J. M., and Saccucci, M. S. (1990), "Exponentially Weighted Moving Average Control Schemes: Properties and Enhancements," *Technometrics*, 32, 1–12.

Mason, R. L., Champ, C. W., Tracy, N. D., Wierda, S. J., and Young, J. C. (1997), "Assessment of Multivariate Process Control Techniques," *Journal of Quality Technology*, 29, 140–143.

Montgomery, D. C. (1997), *Introduction to Statistical Quality Control* (3rd ed.), New York: Wiley.

Montgomery, D. C., and Woodall, W. H. (1997), "A Discussion of Statistically-Based Process Monitoring and Control," *Journal of Quality Technology*, 29, 121–162.

Moustakides, G. V. (1986), "Optimal Stopping Times for Detecting Changes in Distributions," *The Annals of Statistics*, 14, 1379–1387.

Nishina, K. (1992), "A Comparison of Control Charts From the Viewpoint of Change-Point Estimation," *Quality and Reliability Engineering International*, 8, 537–541.

Nishina, K., and Peng-Hsiung, W. (1996), "Performance of CUSUM Charts From the Viewpoint of Change-Point Estimation in the Presence of Autocorrelation," *Quality and Reliability Engineering International*, 12, 3–8.

Page, E. S. (1954), "Continuous Inspection Schemes," *Biometrika*, 41, 100–114.

Pollak, M. (1985), "Optimal Detection of a Change in Distribution," *The Annals of Statistics*, 13, 206–227.

Rao, C. R. (1989), *Statistics and Truth: Putting Chance to Work*, Fairland, MD: International Co-Operative Publishing House.

Reynolds, M. R., Jr., Amin, R. W., and Arnold, J. C. (1990), "CUSUM Charts With Variable Sampling Intervals," *Technometrics*, 32, 371–384.

Reynolds, M. R., Jr., and Stoumbos, Z. G. (1998), "The SPRT Chart for Monitoring a Proportion," *IIE Transactions on Quality and Reliability Engineering*, 30, 545–561.

——— (1999), "A CUSUM Chart for Monitoring a Proportion When Inspecting Continuously," *Journal of Quality Technology*, 31, 87–108.

——— (2000a), "Monitoring the Process Mean and Variance Using Individual Observations and Variable Sampling Intervals," *Journal of Quality Technology*, (to appear).

——— (2000b), "A General Approach to Modeling CUSUM Charts for a Proportion," *IIE Transactions on Quality and Reliability Engineering, Special Issue on Emerging Trends in Quality Engineering*, 32, 515–535.

Ritov, Y. (1990), "Decision Theoretic Optimality of the CUSUM Procedure," *The Annals of Statistics*, 18, 1464–1469.

Roberts, S. W. (1959), "Control Charts Based on Geometric Moving Averages," *Technometrics*, 1, 239–250.

——— (1966), "A Comparison of Some Control Chart Procedures," *Technometrics*, 8, 411–430.

Ryan, T. P. (2000), *Statistical Methods for Quality Improvement* (2nd ed.), New York: Wiley.

Shewhart, W. A. (1931), *Economic Control of Quality of Manufactured Product*, New York: Van Nostrand.

Shiryaev, A. N. (1963), "On Optimum Methods in Quickest Detection Problems," *Theory of Probability and Applications*, 8, 22–46.

Srivastava, M. S., and Wu, Y. (1994), "Dynamic Sampling Plan in the Shiryaev–Roberts Procedure for Detecting a Change in the Drift of Brownian Motion," *The Annals of Statistics*, 22, 805–823.

Stoumbos, Z. G. (1999), "The Detection and Estimation of the Change Point in a Discrete-Time Stochastic System," *Stochastic Analysis and Applications*, 17, 637–649.

Stoumbos, Z. G., and Jones, L. A. (2000), "On the Properties and Design of Individuals Control Charts Based on Simplicial Depth," *Nonlinear Studies*, (to appear).

Stoumbos, Z. G., Jones, L. A., Woodall, W. H., and Reynolds, M. R., Jr. (2000), "On Shewhart-Type Nonparametric Multivariate Control Charts Based on Data Depth," in *Frontiers in Statistical Quality Control 6*,

eds. H.-J. Lenz and P.-Th. Wilrich, Heidelberg, Germany: Springer-Verlag, 208–228.
Stoumbos, Z. G., and Reynolds, M. R., Jr. (1997), "Control Charts Applying a Sequential Test at Fixed Sampling Intervals," *Journal of Quality Technology*, 29, 21–40.
——— (2000a), "The SPRT Control Chart for the Process Mean With Samples Starting at Fixed Times," *Nonlinear Analysis*, (to appear).
——— (2000b), "Robustness to Non-Normality and Autocorrelation of Individual Control Charts," *Journal of Statistical Computation and Simulation, Special Issue on Statistics: Modeling of Processes in Science and Industry*, 66, 145–187.
Tagaras, G. (1998), "A Survey of Recent Developments in the Design of Adaptive Control Charts," *Journal of Quality Technology*, 30, 212–231.
Tsung, F., Shi, J., and Wu, C. F. J. (1999), "Joint Monitoring of PID-Controlled Processes," *Journal of Quality Technology*, 31, 275–285.
Woodall, W. H. (1986), "Weaknesses of the Economic Design of Control Charts" (letter to the editor), *Technometrics*, 28, 408–410.
Woodall, W. H., and Montgomery, D. C. (1999), "Research Issues and Ideas in Statistical Process Control," *Journal of Quality Technology*, 31, 376–386.
Yashchin, E. (1997), "Change-Point Models in Industrial Applications," *Nonlinear Analysis*, 30, 3997–4006.

Statistics in Preclinical Pharmaceutical Research and Development

Bert GUNTER and Dan HOLDER

Although most statisticians and the public at large are familiar with the role of statistics in human clinical drug trials, advances in the basic science and technology of drug research and development (R&D) have created equally challenging and important opportunities in the preclinical arena. Preclinical pharmaceutical research encompasses all aspects of drug discovery and development, from basic research into how cells and organs work and how disease processes disrupt that work to the development, formulation, and manufacture of drugs. The activities that fall under this rubric include biological and biochemical research using in vitro ("test tube"-based) and in vivo (whole animal) experiments; genomics, the study of gene expression in cells, organisms, and populations to determine the molecular biology of disease; proteomics, the study of protein expression patterns to understand how normal and disease processes differ; design, synthesis, and selection of diverse chemical and natural product "libraries" of compounds to screen for desirable biological activity, often via high throughput, "industrialized" drug screening assays; analytical development for drug research and manufacturing; animal testing of drug candidates for efficacy and metabolism and to determine drug toxicity, teratogenicity (fetal and growth effects), and carcinogenicity; development and scale-up of chemical and fermentation drug manufacturing processes; and drug formulation and stability testing. This list is far from complete.

To put these activities into perspective, it can easily cost more than $1 billion and require 10 to 15 years of R&D to bring out a single new drug, of which only the last 2–3 involve the FDA-reviewed human trials with which statisticians and the public are most familiar. So preclinical activities occupy the bulk of the time and scientific effort. The statistics that support this work cover a broad range of statistical methods. Sample sizes can range from longitudinal case-control studies of 10 or fewer animals (although they may produce thousands of data points from continuous monitoring using sophisticated instruments and telemetry) to hundreds of thousands or millions of multivariate records in drug screening and structure searches. All areas of statistics find useful application, but recent opportunities for nonparametric experimental design, linear and nonlinear longitudinal data modeling, high-dimensional exploration and visualization, inference using exact permutation methods and bootstrapping, and pattern recognition, classification, and clustering of large databases are perhaps noteworthy.

Clearly, in a brief survey like this we can highlight only a couple of examples. We have chosen chemometrics and genomics because they provide good examples of the kind of interdisciplinary, data-rich, and nonstandard issues that are increasingly at the forefront of modern pharmaceutical research. But these examples are just the tip of a vast and fascinating iceberg.

1. CHEMOMETRICS

Roughly speaking, chemometrics is the statistics of (analytical) chemistry data, especially spectroscopy data. Physics and chemistry have developed an arsenal of ingenious tools to probe chemical composition and structure. (A nice internet resource for spectroscopy is www.anachem.umu.se/jumpstation.htm.) These techniques can produce (one- and two-dimensional) spectra of exquisite resolution, often with hundreds or thousands of individual peaks. Digitizing translates them to multivariate vectors of that dimensionality. Chemometrics arose because classical multivariate normal statistical methods were inadequate for such data and related matters of calibration and quality control.

One typical application will give the flavor of the issues. Suppose that one has, say, 200 unknown natural chemical extracts from various biological sources that are tested for antibiotic activity against 30 different pathogens. (Many an-

Bert Gunter is Director, Scientific Staff, Merck Research Laboratories, Rahway, NJ 08540 (E-mail: *bert_gunter@merck.com*) and Dan Holder is Associate Director, Scientific Staff, Merck Research Laboratories, West Point, PA 19486 (E-mail: *dan_holder@merck.com*).

Journal of the American Statistical Association
September 2000, Vol. 95, No. 451, Vignettes